Simple Strategies of Agents in an Evolving Auction Model

Ric D. Herbert[1] **and Paul Turton**[2]

[1]Faculty of Science and Information Technology, The University of Newcastle, Ourimbah, NSW, 2258, Australia.
Email:ric.herbert@newcastle.edu.au
[2]Department of Electrical and Communication Engineering, Institut Teknologi Brunei, Brunei.
Email:paul@itb.edu.bn

Keywords: *Auctions, Agent-based Model, Price-setting strategies*

ABSTRACT

Auctions are a common method of buying and selling goods and services. As an area of research auctions have become increasing important and more generally available to the public with the advent of online auctions for the sale of goods. In an auction, sellers ask a price and buyers bid a price. Other than that there are many and varied rules that can make up an auction (Wurman et al. 1998, Parsons et al. 2005). The most common forms are the single-sided and double-sided. In the single-sided there is a single seller and multiple buyers or, in a reverse auction, a single buyer and multiple sellers. Examples of single sided auctions are: English, first price ascending; Dutch first price descending; First Price Sealed Bid; and, Second Price Sealed Bid or Vickery. In double-sided there are multiple buyers and sellers. A common double-sided auction is the continuous double auction where trades are permitted at any time and buyers and sellers can continually update their bids and asks. Continuous double auctions are commonly used in markets to trade stocks and other commodities.

Auctions have been used in areas other than buying and selling. They have been used in the design of control systems for complex processes in an area known as Market Based Control. For example, computational agents representing temperature controllers in a double-blind computerized auction moderated by a central auctioneer has been used to regulate the temperature in a building (Huberman & Clearwater 1995). Using a continuous double auction and multiple markets Lalis et al. (1998) prototype the allocation of servers and applications over computer networks. Reinicke et al. (2005) also use economic principles in resource allocation for grid computers. Yen (2004) use market driven agents with auctions for internet scheduling. Economic or market based methods have been have also been used in the allocation of sensor resources in sensor networks (Vidal 2003, Mainland et al. 2004, Yujie et al. 2006, e.g.). They have also been used in resource allocation in electricity infrastructure (Kok et al. 2005) and telecommunications infrastructure (Gibney et al. 1998).

In recent times auction models for resource allocation have tended to rely more heavily on economics principles, with concepts like Pareto optimality (Mainland et al. 2004) and utility (Vidal 2003). Dash et al. (2006) limited seller capacity as well as fixed and variable costs are included in a continuous double-auction where buyers and sellers continually post their bids and asks on a billboard and a trade occurs whenever the buy is greater than or equal to the ask. They find this leads to an efficient resource allocation.

One issue is what strategies should buyers and sellers use in an auction? In the Santa Fe Auction (Rust et al. 1994), researchers were invited to submit programs that implemented bidding strategies in a double auction. The programs in the tournament shared $10000 in relationship to how successful they were in the tournament. The most successful strategy was a simple strategy.

Other researchers have also found that simple strategies are successful. Gode & Sunder (1993) introduced zero intelligence traders and found them very successful. Cliff (2006) introduces zero intelligence plus traders and finds they outperform human traders in some auctions (Cliff & Bruten 1998, Walia et al. 2003).

In this paper we examine a number of simple strategies for agents in an auction. We allow the agents to evolve with unsuccessful agents dying and being replaced by agents that inherit the most successful strategy. We ask the question as to whether a particular strategy dominates or whether a steady-state of a mixture of strategies evolves.

We find that one strategy is best for buyer agents. However the same strategy is worst for the seller agents.

1 INTRODUCTION

Auctions are a common method of buying and selling goods and services. As an area of research auctions have become increasing important and more generally available to the public with the advent of online auctions for the sale of goods. In an auction, sellers ask a price and buyers bid a price. Other than that there are many and varied rules that can make up an auction (Wurman et al. 1998, Parsons et al. 2005). The most common forms are the single-sided and double-sided. In the single-sided there is a single seller and multiple buyers or, in a reverse auction, a single buyer and multiple sellers. Examples of single sided auctions are: English, first price ascending; Dutch first price descending; First Price Sealed Bid; and, Second Price Sealed Bid or Vickery. In double-sided there are multiple buyers and sellers. A common double-sided auction is the continuous double auction where trades are permitted at any time and buyers and sellers can continually update their bids and asks. Continuous double auctions are commonly used in markets to trade stocks and other commodities.

Auctions have been used in areas other than buying and selling. They have been used in the design of control systems for complex processes in an area known as Market Based Control. For example, computational agents representing temperature controllers in a double-blind computerized auction moderated by a central auctioneer has been used to regulate the temperature in a building (Huberman & Clearwater 1995). Using a continuous double auction and multiple markets Lalis et al. (1998) prototype the allocation of servers and applications over computer networks. Reinicke et al. (2005) also use economic principles in resource allocation for grid computers. Yen (2004) use market driven agents with auctions for internet scheduling. Economic or market based methods have been have also been used in the allocation of sensor resources in sensor networks (Vidal 2003, Mainland et al. 2004, Yujie et al. 2006, e.g.). They have also been used in resource allocation in electricity infrastructure (Kok et al. 2005) and telecommunications infrastructure (Gibney et al. 1998).

In recent times auction models for resource allocation have tended to rely more heavily on economics principles, with concepts like Pareto optimality (Mainland et al. 2004) and utility (Vidal 2003). Dash et al. (2006) limited seller capacity as well as fixed and variable costs are included in a continuous double-auction where buyers and sellers continually post their bids and asks on a billboard and a trade occurs whenever the buy is greater than or equal to the ask. They find this leads to an efficient resource allocation.

One issue is what strategies should buyers and sellers use in an auction? In the Santa Fe Auction (Rust et al. 1994), researchers were invited to submit programs that implemented bidding strategies in a double auction. The programs in the tournament shared \$10000 in relationship to how successful they were in the tournament. The most successful strategy was a simple strategy.

Other researchers have also found that simple strategies are successful. Gode & Sunder (1993) introduced zero intelligence traders and found them very successful. Cliff (2006) introduces zero intelligence plus traders and finds they outperform human traders in some auctions (Cliff & Bruten 1998, Walia et al. 2003).

In this paper we examine a number of simple strategies for agents in an auction. We allow the agents to evolve with unsuccessful agents dying and being replaced by agents that inherit the most successful strategy. We ask the question as to whether a particular strategy dominates or whether a steady-state of a mixture of strategies evolve.

2 THE AUCTION MODEL

The auction model we use is a continuous double auction where buyers and sellers post their bids on a billboard in each round and a trade occurs whenever a bid is greater than or equal to an ask. The model is considered as a game which is repeated over time.

2.1 Key Features

The main characteristics of the model are:

- Auction environment:
  - It is a double auction with multiple buyer agents bidding and multiple supplier agents independently asking. That is, there is a set of n buyers, $B = \{b_1, b_2, ..., b_n\}$; and m sellers, $S = \{s_1, s_2, ..., s_m\}$.
  - The auction is continuous in that buyers and sellers are free to change their bids, but for convenience we run the auction more in the manner of a game in rounds of a short time period where each buyer and seller makes a decision for the next round. We donate rounds by k.
  - The asks and bids are posted on a billboard.
  - Agents are restricted to a single type in that agents are either buyers or sellers. Hence agents cannot buy and then resell.

- Agents buy and sell a single homogenous good, g, of unlimited quantity. That is, for this paper, we ignore supply side issues.
- The only decision for each agent in each round is to determine it's ask or bid. Asks are denoted $p_{s,i}$ and bids as $p_{b,i}$ for the ith agent.
- The agents have no information about how other agents are bidding or asking, during the current round. Hence the auction could be considered a sealed-bid auction for the current round.
- Depending on their strategy, agents learn the market price after the auction occurs and may use that information as part of their strategy for the next round. The market price, p_m, is the average price of all trades in the round.
- Again depending on their strategy, some agents have a fixed-length memory and can remember price information. In this paper we limit the agent's price information to be the market price and the agent's bid. Thus the ith agent has a memory matrix which store the history, H_i, containing a column of market price and another for it's bid. The memory length determines the number of rows and this may vary for each agent.

• Auctioneer Agent:
  - There is an auctioneer agent or monitor, $\mathcal{M}$, who clears the market after each round.
  - The monitor randomly examines pairs of random asks and random bids and if the bid is greater than or equal to the ask then a trade occurs. That is, a trade occurs giving a trading price of $p_l^t = p_{b,i}$ if $p_{b,i} \geq p_{s,j}$ making the lth trade between the ith buyer and jth seller. The trades for a set for the round, $\mathcal{T} = \{p_{k,1}^t, p_{k,2}^t, ...p_{k,o}^t\}$, where o is the number of successful trades in this round.
  - Unmatched bids and asks are not cleared, so that not all agents buy or sell.

• Buyer Agents:
  - There are a fixed number of buyer agents, n.
  - Each buyer has a stock of funds, f_i, from which it purchases goods.
  - Buyers receive an income, y_i, at each round which adds to their stock of funds.
  - Each buyer has a single strategy, σ_i, which it uses in determining the next bid.
  - The objective of a buyer is to accumulate funds. Buyer are wealth maximisers.

• Seller Agents:
  - There are a fixed number of seller agents, m. The number of sellers is possibly different number of buyers ($n \neq m$).
  - Sellers make a profit by selling goods, which adds to their stock of funds, f_i.
  - The sellers profit from a trade is the price they obtain minus a fixed cost.
  - Each seller has a single strategy, σ_i, which it uses in determining the next bid.
  - Sellers receive an income, y_i, at each round which adds to their stock of funds. In this paper this is zero.
  - The objective of sellers is also to accumulate funds.

Notice that the buyer and seller agents have identical attributes, and can be modeled by the same software objects.

2.2 Agent Strategies

In this paper we focus on three simple agent strategies, and these define three types of traders. Both buyers and seller use the same set of strategies. The traders are:

• Zero Intelligence Traders (ZIT). These traders simple use a random price for their bid or ask.

• Market Price Traders (MPT). These agents use the average market price from the last round of the auction.

• Historical Price Traders (HPT). Each agent has a memory of their last η trades, where η is a parameter of the model and may be different for buyer (η_b) and seller (η_s) agents. Historical Price Traders take their average price over their memory length of their bids or asks and use that price in the next round.

Hence $\sigma_i \in \{ZIT, MPT, HPT\}$. Agents with the first two strategies are memory-free agents whereas the third agent is a memory based agent (Bagnall & Toft 2006). Other types of strategies can be found in He et al. (2003).

2.3 Evolution

In this paper we introduce a simple evolutionary structure for agents within the model. The basis of this evolutionary mechanism is that successful agents 'retire' from the game, and unsuccessful agents exit from the auction ('die'). Success for sellers and buyers is defined as funds being above some critical level, $f_i > \overline{f}$. Unsuccessful agents are those whose funds are below some critical level, $f_i < \underline{f}$.

We keep the numbers of buyer agents and the number of seller agents constant during the game. When a successful agent retires it is replaced by a new agent who is born with random initial values $(m_0, f_0, \overline{f_0}, \underline{f_0}, y_0)$ except that it retains the same strategy as the agent it replaces. When an unsuccessful agent dies, a new agent of the same type is born. This agent has random initial values except that it inherits the most successful strategy of agents of its type. The most successful strategy is the strategy of the agent with the most funds.

Evolution occurs after each round in the auction. Evolution can be described as a software method on the set of buyers and sellers.

2.4 The Auction Model

Our auction model can be considered as a game, G, but we do not examine it here with game theory. It is played in rounds with a time limit of $\overline{k}$ rounds. A **descriptor** of the auction game is:

Auction Game: $G =< S, B, g, \overline{k} >$.

Monitor Agent: $\mathcal{M} =< \mathcal{T}, k >$.

Seller Agent: $s =< p, f, y, \sigma, \overline{f}, \underline{f} >$.

Buyer Agent: $b =< p, f, y, \sigma, \overline{f}, \underline{f} >$.

Evolution: $\mathcal{E} =< k, B, S >$.

A **protocol** for the auction is:

1. Initialise G.
2. Round:
   (a) Increment round: $k \leftarrow k + 1$.
   (b) Determine agent bid: p_b, p_s.
   (c) Calculate trades $\mathcal{T}$.
   (d) Calculate market statistics.
   (e) Evolution: $\mathcal{E}$.
3. Finalise:

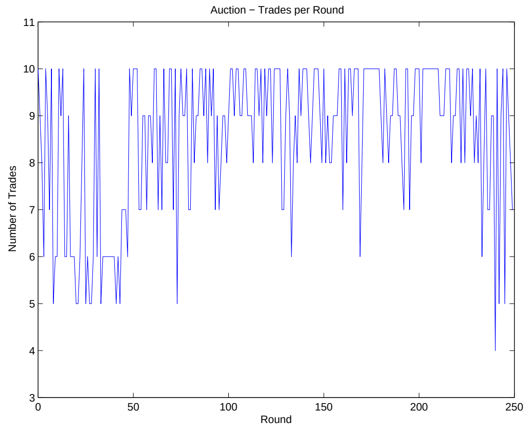

Figure 1. Number of trades over time.

3 RESULTS AND DISCUSSION

We start by initialising the model with 10 sellers and 15 buyers in a game of 250 rounds ($n = 15, m = 10, \overline{k} = 250$). We normalise prices to be between 0 and 1. The historical price traders have a memory length of 3 ($m = 3$). Figure 1 presents the number of trades per round. A trade takes place when the randomly matched buyer agent has a bid price greater than or equal to the asking price of the randomly picked seller agent. The maximum possible number of trades per round is 10, but from the Figure it can be seen that this only occurs in about 40% of rounds. The mean number of trades per round is 8.6 and the variance is 2.5, indicating that most sellers complete a trade in each round. The Figure does not show any significant overall pattern as the strategies of the players in the auction game evolve over time.

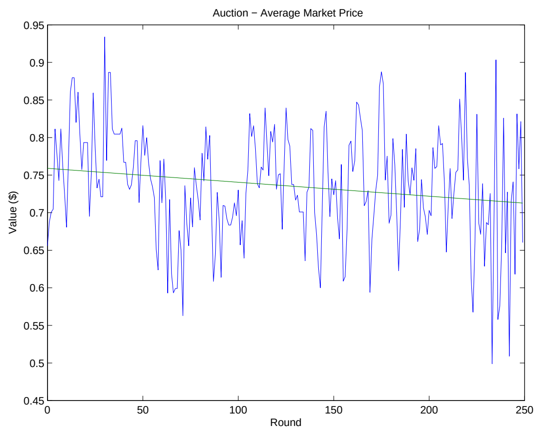

Figure 2. Average market price over time.

Figure 2 presents the average market price for the auction. The figure shows that there is considerable price volatility. The mean average market price for this auction is 0.74 and the variance is 0.01. The figure also plots the linear (least squares) trend line for the average market price. The trend shows a gradually decreasing price over the rounds in the auction game.

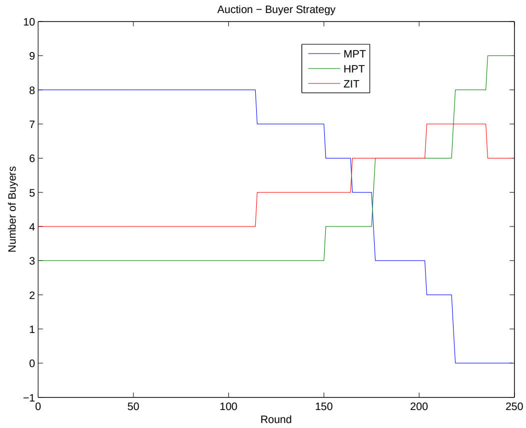

Figure 3. Number of buyers by strategy over time.

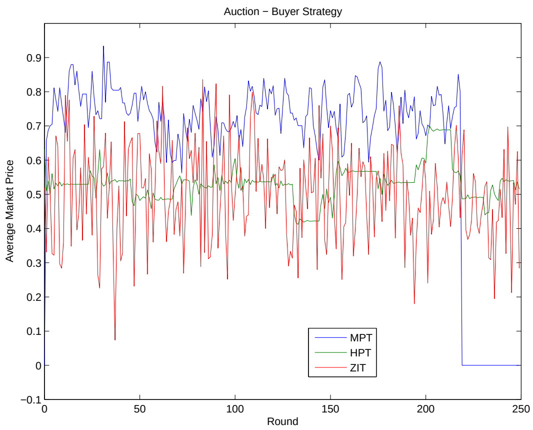

Figure 4. Average price paid by buyers by strategy over time.

We now turn to consider the evolution of strategies within the game. Initially buyer and seller agents are randomly allocated a pricing strategy from the available set. Over time evolution occurs with agents. Figure 3 plots the evolution of buyer strategies over time. The figure shows that the MPT strategy buyers decline in number very quickly and eventually are completely replaced by HPT and ZIT strategy buyers. HPT strategy buyers start the game in the minority and are in the majority by the end of the game. This strategy is evolutionary the most successful.

Why is the MPT strategy so unsuccessful for buyers? Figure 4 shows the average bid by buyers of different strategy types. MPT buyers, on average, bid a higher price than the other strategy types. This may explain their lack of success.

With seller agents the evolutionary effects also occur quickly. This can be seen from Figure 5, where

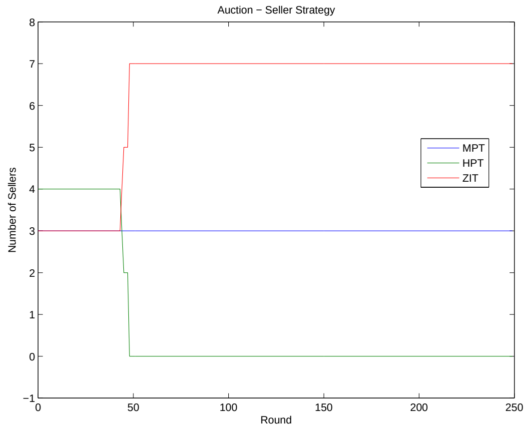

Figure 5. Number of sellers by strategy over time.

the number of seller agents of each strategy type are plotted over time. The figure shows that the evolution occurs quickly and settles into a new and maintained steady-state. HPT seller agents rapidly die out and are replaced by ZIT sellers. MPT strategy seller agents remain constant in number in this example. Figure 6 shows the ask of sellers. Again, the highest ask is for the group that die out. In this case it is the HPT strategy that is unsuccessful.

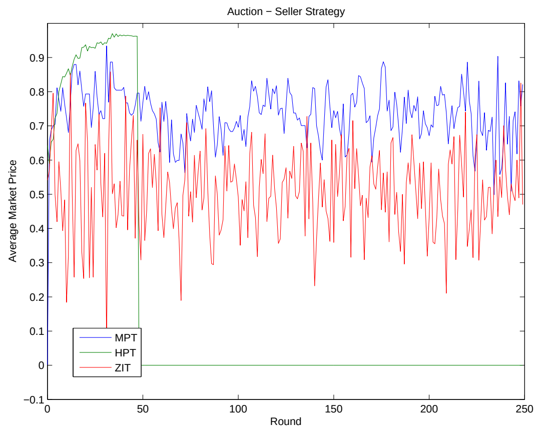

Figure 6. Average price received by sellers by strategy over time.

Notice that with sellers the transient dynamics of evolution are fast. The evolved steady-state is reached before there are evolutionary changes in the number of buyer agents.

We now turn to the question of whether or not this result is simply a product of the particular auction game or is more general in nature. In an attempt to answer this question the model was run a number of times with each run giving us different random variables for buyer and seller strategies as well as different starting values for all other variables. The only constant values to aid in comparisons are number of buyers and sellers and the time horizon. The results

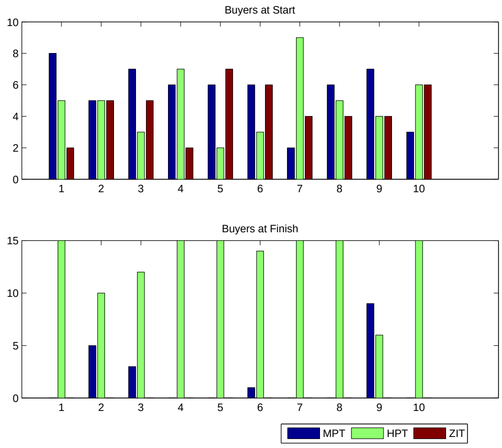

Figure 7. Start and finishing number of buyers by strategy.

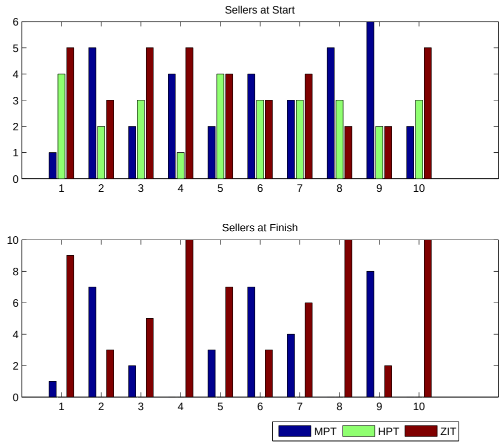

Figure 8. Start and finishing number of sellers by strategy.

are shown in Figure 7 and Figure 8. It can be seen that whilst the number of buyers and sellers of each strategy vary at the start of the game at the end of the game two distinct patterns emerge: – for buyers, HPT buyers always win and for sellers, HPT sellers always lose. We define win to mean they are in the majority and loose that there are the least number if any of them.

A deal results in two successes, a successful buyer and a successful seller. In order to explain why the HPT buyers are so successful it may be helpful to consider the successful sellers. With the sellers it is not as clear-cut as with the buyers, there is no one dominant winner. The MPT sellers are more successful than the ZIT sellers (the HPT sellers having no success at all!). From Figure 8 it can be seen that the MPT sellers won 70% of the games, the ZIT sellers won the other 30% of the games. So the MPT sellers are almost twice as successful as the ZIT sellers? The MPT sellers use the average market price of the previous round as their ask price, so why is this good? The HPT buyers derive their bid price from their memory length - which looking at their dominance in the later stages of the game, would come from very recent game history. This would allow some price stability on the buyer side. On the seller side there is a reasonable amount of instability provided by the ZIT sellers, they will generate low asks that will fuel the deals with the HPT buyers, this will have the effect of providing the HPT buyers with cheap deals. The MPT sellers will again provide some price stability and average out any large swings.

While we may expect sellers and buyers to have opposite results on the basis of a non-zero sum game, this game allows for sellers to do better as buyers may offer more and for buyers to do better as sellers may ask less. So it is possible for both buyers and sellers to do well in this game. HPT sellers however do not seem to gain the benefit of this loosing all the time. Having a knowledge of their own past market performance is of no use in this model. They are not rewarded with future sales - this seems to be counter-intuitive and may be to do with the assumptions inherent in the model. If HPT buyers perform well shouldn't HPT sellers perform as well?

4 CONCLUSION

In this paper we have consider the evolution of agent's strategies in a model of an auction. There were seller agents and buyer agents. The auction model we use is a continuous double auction where buyers and sellers post their bids on a billboard in each round and a trade occurs whenever a bid is greater than or equal to an ask. Within this model we considered the issue of a variety of agent strategies. For both types of agents we consider the same strategies. Some agent's bid or ask was determined randomly (Zero Intelligence Traders); others determined their bid or ask by simply mimicking the mean market price (Market Price Traders); and the final strategy was for agent's to use the mean of their fixed length memory of previous bids or asks (Historical Price Traders).

The model had an evolutionary mechanism where successful agents reproduced and unsuccessful agents died. We found that the historical price strategy was best for buyer agents. But the same strategy was worst for the seller agents. We speculate that this may be because this stored knowledge of the sellers is crowded out by the random behavior of the ZIT sellers, they do not care about their price and so slant the market providing cheap products for the HPT buyers at the expense of the HPT sellers.

In future work we propose to widen the type and number of games to see if the results are more general. What happens to the results as we vary the number of agents of each type and modify other important parameters? We also plan to introduce other strategies and allow buyer and seller agents to choose from a variety of strategies.

References

Bagnall, A. & Toft, I. (2006), 'Autonomous adaptive agents for single seller sealed bid auctions', *Autonomous Agents and Multi-Agent Systems* **12**(3), 259–292.

Cliff, D. (2006), Evolutionary optimization of zip60: a controlled explosion in hyperspace, *in* 'GECCO '06: Proceedings of the 8th annual conference on Genetic and evolutionary computation', ACM Press, New York, NY, USA, pp. 1621–1628.

Cliff, D. & Bruten, J. (1998), Simple bargaining agents for decentralized market-based control, *in* 'Proceedings of the 12th European Simulation Multiconference on Simulation - Past, Present and Future', SCS Europe, pp. 478–485.

Dash, R. K., Vytelingum, P., Rogers, A., David, E. & Jennings, N. R. (2006), 'Market-based task allocation mechanisms for limited capacity suppliers', *IEEE Trans on SMC* . forthcoming.

Gibney, M. A., Jennings, N. R., Vriend, N. J. & Griffiths, J. M. (1998), *Market-based call routing in telecommunications networks using adaptive pricing and real bidding*, Vol. 1437 of *Lecture Notes in Computer Science*, Springer, Berlin, chapter 3, pp. 46–61.

Gode, D. K. & Sunder, S. (1993), 'Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality', *The Journal of Political Economy* **101**(1), 119–137.

He, M., fung Leung, H. & Jennings, N. R. (2003), 'A fuzzy-logic based bidding strategy for autonomous agents in continuous double auctions', *IEEE Transactions on Knowledge and Data Engineering* **15**(6), 1345–1363.

Huberman, B. & Clearwater, S. H. (1995), A multi-agent system for controlling building environments, *in* V. Lesser, ed., 'Proceedings of the 1st International Conference on Multi-agent Systems', pp. 171–176.

Kok, J. K., Warmer, C. J. & Kamphuis, I. G. (2005), Powermatcher: Multiagent control in the electricity infrastructure, *in* 'AAMAS'05: Proceedings', ACM Press, New York, NY, USA, pp. 75–82.

Lalis, S., Nikalaou, C., Papadakis, D. & Marazakis, M. (1998), Market-driven service allocation in a qos-capable environment, *in* 'ICE '98: Proceedings of the first international conference on Information and computation economies', ACM Press, New York, NY, USA, pp. 92–100.

Mainland, G., Kang, L., Lahaie, S., Parkes, D. C. & Welsh, M. (2004), Using virtual markets to program global behavior in sensor networks, *in* 'EW11: Proceedings of the 11th workshop on ACM SIGOPS European workshop: beyond the PC', ACM Press, New York, NY, USA, pp. 1–6.

Parsons, S., Marcinkiewicz, M. & Niu, J. (2005), Everything you wanted to know about double auctions, but were afraid to (bid or) ask, Working paper, City University of New York.
URL: *http://www.sci.brooklyn.cuny.edu/ parsons-/projects/mech-design/publications/cda.pdf*

Reinicke, M., Streitberger, W. & Eymann, T. (2005), 'Evaluation of service selection techniques in service oriented computing networks', *Multiagent Grid Syst.* **1**(4), 271–285.

Rust, J., Miller, J. & Palmer, R. (1994), 'Characterizing effective trading strategies – insights from a computerized double auction tournament', *Journal of Economic Dynamics and Control* **18**, 61–96.

Vidal, J. M. (2003), 'A method for solving distributed service allocation problems', *Web Intelli. and Agent Sys.* **1**(2), 139–146.

Walia, V., Byde, A. & Cliff, D. (2003), Evolving market design in zero-intelligence trader markets, *in* 'CEC 2003. IEEE International Conference on E-Commerce', pp. 157–164.

Wurman, P. R., Wellman, M. P. & Walsh, W. E. (1998), The michigan internet auctionbot: a configurable auction server for human and software agents, *in* 'AGENTS '98: Proceedings of the second international conference on Autonomous agents', ACM Press, New York, NY, USA, pp. 301–308.

Yen, B.P.-C.and Wu, O. (2004), 'Internet scheduling environment with market-driven agents', *IEEE Transactions on Systems, Man and Cybernetics, Part A* **34**(2), 281–289.

Yujie, W. Z. Y., Peng, J. & Shu, W. (2006), Market-based adaptive task scheduling for sensor networks, *in* 'WiCOM 2006: Proceedings of Wireless Communications, Networking and Mobile Computing', IEEE, New York, NY, USA, pp. 1–5.